

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.

REF / HP

06/21/2004 15:02 FAX 7132388000

CONLEY, ROSE

RECEIVED
CENTRAL FAX CENTER

JUN 21 2004

008

OFFICIAL

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Applicants:	Shubhendu S. Mukherjee et al.	Confirmation No.:	2964
Serial No.:	09/652,322	Group Art Unit:	2141
Filed:	08/31/2000	Examiner:	Q. N. Nguyen
For:	Priority Rules For Reducing Network Message Routing Latency	Docket No.:	200301777-1

PRELIMINARY AMENDMENT

Mail Stop RCE
Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

Date: June 21, 2004

Sir:

In response to the final Office action of April 7, 2004, please enter this Preliminary Amendment and corresponding Request for Continued Examination ("RCE") and amend the above-identified application as follows:

Amendments to the Claims are reflected in the listing of claims which begins on page 2 of this paper.

Remarks/Arguments begin on page 13 of this paper.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (Currently amended) A distributed multiprocessing computer system, comprising:

a plurality of microprocessor units coupled to each other, wherein each microprocessor unit comprises;

a router to route message packets between said microprocessor units, and wherein said router prioritizes message packets based upon type of message packet, age of the message packet, size of the message packet and source of the message packet;

a plurality of network input ports and network output ports connecting said plurality of microprocessor units to form a computer network, wherein each of said network input ports couples to one or more associated local arbiters in the router, each of said local arbiters operable to select a message packet among message packets waiting at the network input port.

2. (Previously presented) The distributed multiprocessing computer system of claim 1 wherein said router includes a plurality of timers that indicate when a message packet must be immediately dispatched.

3. (Previously presented) The distributed multiprocessing computer system of claim 1 wherein said microprocessor unit further includes a plurality of microprocessor input ports and microprocessor output ports that allow the

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

exchange of message packets between hardware functional units in the microprocessor unit and between microprocessor units.

4. (Original) The distributed multiprocessing computer system of claim 3 wherein each of said microprocessor input ports couples to local arbiters in the router, each of said local arbiters able to select a message packet among message packets waiting at the microprocessor input port.

5. (Original) The distributed multiprocessing computer system of claim 4 wherein each of said network output ports and microprocessor output ports couples to a global arbiter in the router that selects a message packet from message packets nominated by the local arbiters of said network input ports and microprocessor input ports.

6. (Original) The computer system of claim 5 wherein if a first message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the first message packet type from the global arbiter of the destination network output port or microprocessor output port.

7. (Original) The computer system of claim 6 wherein if a second message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the second message packet type from the global arbiter of the destination network output port or microprocessor output port.

8. (Original) The computer system of claim 7 wherein if a third message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the third message packet type from the global arbiter of the destination network output port or microprocessor output port.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

9. (Original) The computer system of claim 8 wherein if a fourth message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the fourth message packet type from the global arbiter of the destination network output port or microprocessor output port.

10. (Original) The computer system of claim 9 wherein if a fifth message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the fifth message packet type from the global arbiter of the destination network output port or microprocessor output port.

11. (Original) The computer system of claim 10 wherein if a sixth message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the sixth message packet type from the global arbiter of the destination network output port or microprocessor output port.

12. (Original) The computer system of claim 11 wherein if a seventh message packet type is ready to be dispatched from the network input port or microprocessor input port, the local arbiter requests service for the seventh message packet type from the global arbiter of the destination network output port or microprocessor output port.

13. (Original) The computer system of claim 5 wherein said network output port global arbiter or microprocessor output port global arbiter selects said message packet Least-Recently-Granted from the network input ports, then Least-Recently-Granted from the microprocessor input ports if said network output port or microprocessor output port is idle.

Appl. No.: 09/652,322
Amtd. dated June 21, 2004
Reply to Office action of April 7, 2004

(A) 14. (Currently amended) A method of routing messages in a distributed multiprocessor computer system, comprising:

selecting a message packet at each of a plurality of microprocessor router input ports from message packets buffered at each input port based on the type of message packet; and

transmitting from an idle microprocessor router output port a message packet chosen from the plurality of selected microprocessor router input port message packets, said message packet chosen for transmission by the output port based on the microprocessor router input port priority,

wherein said selecting and said transmitting reduces routing latency of the distributed multiprocessor computer system by assigning message packets already in transit higher priority than new message packets and by implementing a combination of type-based and least-recently-granted routing algorithms for both message packets already in transit and new message packets.

15. (Original) The method of claim 14 wherein said selecting a message packet includes the step of:

determining if a Block Response packet is ready to be dispatched from the input port buffer; and

if the Block Response packet is ready, selecting the Block Response packet.

16. (Original) The method of claim 15 wherein said selecting a message packet includes the step of:

if no Block Response packet is ready, determining if a Acknowledgment packet is ready to be dispatched from the input port buffer; and

if the Acknowledgment packet is ready, selecting the Acknowledgment packet.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

17. (Original) The method of claim 16 wherein said selecting a message packet includes the step of:

if no Acknowledgment packet is ready, determining if an Invalidation Broadcast packet is ready to be dispatched from the input port buffer; and
if the Invalidation Broadcast packet is ready, selecting the Invalidation Broadcast packet.

18. (Original) The method of claim 17 wherein said selecting a message packet includes the step of:

if no Invalidation Broadcast packet is ready, determining if a Forward packet is ready to be dispatched from the input port buffer; and
if the Forward packet is ready, selecting the Forward packet.

19. (Original) The method of claim 18 wherein said selecting a message packet includes the step of:

if no Forward packet is ready, determining if a Request packet is ready to be dispatched from the input port buffer; and
if the Request packet is ready, selecting the Request packet.

20. (Original) The method of claim 19 wherein said selecting a message packet includes the step of:

if no Request packet is ready, determining if a Write I/O packet is ready to be dispatched from the input port buffer; and
if the Write I/O packet is ready, selecting the Write I/O packet.

21. (Original) The method of claim 20 wherein said selecting a message packet includes the step of:

if no Write I/O packet is ready, determining if a Read I/O packet type is ready to be dispatched from the input port buffer; and
if the Read I/O packet is ready, selecting the Read I/O packet.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

22. (Original) The method of claim 14 wherein said transmitting a message packet includes the step of prioritizing said message packet Least-Recently-Granted from network input ports, then Least-Recently-Granted from microprocessor input ports, wherein said network input ports and said microprocessor input ports are microprocessor router input ports.

23. (Currently amended) A distributed multiprocessing computer system, comprising:

means for selecting a message packet at each of a plurality of microprocessor router input ports from message packets buffered at each input port based on the type of message packet wherein message packets that pass through at least three microprocessors to reach a destination are given highest priority; and

means for transmitting from an idle microprocessor router output port a message packet chosen from the plurality of selected microprocessor router input port message packets, said message packet chosen for transmission by the output port based on the microprocessor router input port priority,

wherein operation of said means for selecting and said means for transmitting reduces routing latency of the distributed multiprocessing computer system.

24. (Currently amended) A distributed multiprocessing computer system, with a plurality of microprocessors, which comprise:

a router to route message packets between said microprocessors, and wherein said router prioritizes message packets based upon type of message packet, age of the message packet, and source of the message packet;

a plurality of network input ports and network output ports connecting said plurality of microprocessors to form a computer network, wherein each of said network input ports couples to one or more associated

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

local arbiters in the router, each of said local arbiters operable to select a message packet among message packets waiting at the network input port;

~~wherein each of said microprocessors further includes a plurality of microprocessor input ports and microprocessor output ports that allow the exchange of message packets between hardware functional units in the microprocessor and between microprocessors,~~

wherein each network input port and each microprocessor input port is assigned a different priority and wherein all network input ports have higher priority than the microprocessor input ports; and

a disk drive coupled to each of said plurality of microprocessors.

25. (Previously presented) The distributed multiprocessing computer system of claim 23 wherein the message packets are selected in an order based on the type of message packet and wherein the order from highest priority to lowest priority consists of Block Response packet type, Acknowledgement packet type, Invalidation Broadcast packet type, Forward packet type, Request packet type, Write I/O packet type, and Read I/O packet type.

26. (Previously presented) The distributed multiprocessing computer system of claim 24 wherein said router prioritizes message packets based on the source of the message packet comprises the router giving message packets associated with network input ports higher priority than message packets associated with microprocessor input ports.

27. (Previously presented) The distributed multiprocessing computer system of claim 26 wherein said router prioritizes message packets based on the source of the message packet further comprises the router giving message packets associated with I/O ports lowest priority and wherein the I/O ports are associated with the disk drive coupled to each of said plurality of microprocessors.

Appl. No.: 09/652,322
Armdt. dated June 21, 2004
Reply to Office action of April 7, 2004

28. (Previously presented) The distributed multiprocessing computer system of claim 2 wherein the plurality of timers comprise starvation timers and drain timers.

29. (Previously presented) The distributed multiprocessing computer system of claim 5 wherein the global arbiter selects and outputs the message packet from message packets based on an input port hierarchy.

30. (New) A system, comprising:
a plurality of microprocessors coupled to each other and configured to prioritize message packets transferred between two or more of the microprocessors and message packets transferred between two or more hardware units within each of the microprocessors,
wherein each of the microprocessors comprises,
network ports configured to receive message packets from other microprocessors and output message packets to other microprocessors;
microprocessor ports configured to receive message packets from the hardware units and output message packets to the hardware units; and
a router coupled to the network ports and the microprocessor ports,
wherein the router assigns the network ports higher priority than the microprocessor ports and prioritizes message packets associated with each of the network ports and each of the microprocessor ports according to a predetermined prioritization and according to a least-recently-granted prioritization.

31. (New) The system of claim 30 wherein the predetermined prioritization is based on message packet types.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

32. (New) The system of claim 31 wherein the message packet types are selected from a group of packets types consisting of: block response packets, acknowledgement packets, invalidation broadcast packets, forward packets, request packets, write I/O packets and read I/O packets.

33. (New) The system of claim 32 wherein block response packets have highest priority.

34. (New) The system of claim 31 wherein the predetermined prioritization is based on a source of the message packet.

35. (New) The system of claim 34 wherein the source of the message packet is selected from a group of sources consisting of: a first network port, a second network port, a third network port, a fourth network port, a cache control unit, a memory controller and an input/output port.

36. (New) The system of claim 35 wherein the first network port, the second network port, the third network port, the fourth network port, the cache control unit, the memory controller and the input/output port are each assigned a different predetermined priority and wherein the first, second, third and fourth network ports are assigned a higher priority than the cache control unit, the memory controller and the input/output port.

37. (New) The system of claim 36 wherein the cache control unit and the memory controller are assigned a higher priority than the input/output port.

38. (New) The system of claim 30 wherein the router further comprises:
local arbiters;
global arbiters; and
input buffers,

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

wherein the local and global arbiters implement algorithms to carry out the least-recently-granted prioritization.

39. (New) The system of claim 38 wherein the algorithms detect when a message packet is denied transmission for more than a threshold amount of time due to at least one of a type of the message packet and a source of the message packet.

40. (New) The system of claim 38 wherein the algorithms detect when a message packet is denied transmission for more than a threshold amount of time due to the local arbiters and the global arbiters operating independently of each other.

41. (New) The system of claim 38 wherein the router is configured to periodically change the value of a bit associated with a message packet to identify the message packet as non-current.

42. (New) The system of claim 41 wherein the router further comprises a starvation timer that counts down from a maximum value to zero after a message packet is identified as non-current.

43. (New) The system of claim 42 wherein the router further comprises buffers reserved to receive message packets that cause the starvation timer to reach zero.

44. (New) The system of claim 43 wherein the router further comprises a drain timer that is activated when the starvation timer reaches zero and when a target buffer of a message packet associated with an expired starvation timer is available.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

45. (New) The system of claim 44 wherein, when the drain timer expires, the router dispatches only message packets that are non-current until there are no message packets associated with an expired starvation timer.

46. (New) The system of claim 38 wherein the local and global arbiters are configured to perform tests before transmitting a particular message packet, wherein the tests are selected from a group of tests consisting of: checking if a network can accept the particular message packet, checking whether the particular message packet is selected by more than one local arbiter, checking a message packet type, checking if an output port of the router is busy, checking if a destination can accept the particular message packet and checking if a starved message packet is waiting at the input buffers.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

REMARKS/ARGUMENTS

Applicants received the final Office Action dated April 7, 2004, in which the Examiner: (1) rejected claims 1-5, 13-14, 22-24 and 26-29 as obvious in view of U.S. Patent No. 5,970,232 ("Passint") and U.S. Patent No. 5,987,518 ("Gotwald"); and (2) rejected claims 6-12 and 15-21 as obvious in view of Passint, Gotwald and U.S. Patent No. 6,282,195 ("Miller"). In this Response, Applicants amend claims 1, 14, 23 and 24. Applicants also add claims 30-46. Claims 1-46 are pending. Based on the arguments and amendments contained herein, Applicants respectfully request reconsideration and allowance of the pending claims.

I. CLAIM REJECTIONS

A. CLAIM 1

Amended claim 1, in part, requires a "router [that] prioritizes message packets based upon...size of the message packet" (supported, at least, on page 25, lines 4-9, of Applicants' specification). None of the references cited by the Examiner teaches or suggests this limitation. Specifically, Passint teaches that the router chips 50 do not assume any particular message length (see col. 11, lines 23-25) and, therefore, does not teach or suggest "a router" that "prioritizes message packets based upon size of the message packet" as required in claim 1. Further, Gotwald teaches that priority can be based on source address, destination address, the data type and the connection type (see col. 4, lines 61-64). However, Gotwald does not teach or suggest a router that "prioritizes message packets based upon size of the message packet" as required in claim 1. None of the references cited by the Examiner, nor combinations of the references, teaches or suggests the above limitation. For at least this reason, Applicants submit that claims and all claims that depend from claim 1 are allowable.

B. CLAIM 14 (Allowed)

Amended claim 14, in part requires "reduc[ing] routing latency...by assigning message packets already in transit higher priority than new message packets and by implementing a combination of type-based and least-recently-granted routing algorithms for both message packets already in transit and new

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

message packets" (supported, at least, on page 26, lines 1-18, of Applicants' specification). None of references cited by the Examiner teaches or suggest this limitation.

While Passint teaches prioritizing messages based on age and header information (see col. 11, lines 27-30 and 56-58), Passint does not teach or suggest that header information and age are related to prioritization based on "message packets already [being] in transit" as required in claim 14. Also, while Gotwald teaches that priority can be based on source address, destination address, the data type and the connection type, Gotwald does not teach or suggest "assigning message packets already in transit higher priority than new message packets" as required in claim 14.

Further, neither Passint nor Gotwald teaches or suggest "implementing a combination of type-based and least-recently-granted routing algorithms for both message packets already in transit and new message packets" as required in claim 14. The Examiner recognizes that Passint does not teach or suggest prioritizing message packets based on type (see Office Action, page 3, last paragraph) and cites Gotwald as teaching prioritizing based on message type. However, Gotwald only teaches that priority can be based on source address, destination address, the data type and the connection type. Gotwald does not teach or suggest "implementing a combination of type-based and least-recently-granted routing algorithms for both message packets already in transit and new message packets" as required in claim 14. None of the references, nor combinations of the references, cited by the Examiner teaches or suggests the above limitations. For at least these reasons, Applicants submit that claim 14 and all claims that depend from claim 14 are allowable.

C. CLAIM 23 (Allowed)

Amended claim 23, in part, requires that "message packets that pass through at least three microprocessors to reach a destination are given highest priority" (supported, at least, on page 5, lines 6-8 and page 25, lines 4-9, of Applicants' specification). None of the references cited by the Examiner teaches or suggests this limitation.

Appl. No.: 09/652,322
Amdt. dated June 21, 2004
Reply to Office action of April 7, 2004

As described previously, Passint teaches basing priority on header information and age, while Gotwald teaches basing priority on source address, destination address, the data type and the connection type. However, neither Passint nor Gotwald, nor a combination of Passint and Gotwald, teaches or suggests that "message packets that pass through at least three microprocessors to reach a destination are given highest priority" as required in claim 23. For at least these reasons, Applicants submit that claim 23 and all claims that depend from claim 23 are allowable.

D. CLAIM 24 (Allowed)

Amended claim 24, in part, requires that "each network input port and each microprocessor input port is assigned a different priority and wherein all network input ports have higher priority than the microprocessor input ports" (supported, at least, on page 26, lines 1-18, of Applicants' specification). None of the references cited by the Examiner teaches or suggests these limitations.

As described previously, Passint teaches basing priority on header information and age, while Gotwald teaches basing priority on source address, destination address, the data type and the connection type. However, neither Passint nor Gotwald, nor a combination of Passint and Gotwald, teaches or suggests that "each network input port and each microprocessor input port is assigned a different priority and wherein all network input ports have higher priority than the microprocessor input ports" as required in claim 24. For at least these reasons, Applicants submit that claim 24 and all claims that depend from claim 24 are allowable.

II. NEW CLAIMS

Claim 30, in part, requires "a router" that "assigns the network ports higher priority than the microprocessor ports and prioritizes message packets associated with each of the network ports and each of the microprocessor ports according to a predetermined prioritization and according to a least-recently-granted prioritization" (supported, at least, on page 16, lines 1-18). None of the references cited by the Examiner teaches or suggests these limitations.

Appl. No.: 09/652,322
Amtd. dated June 21, 2004
Reply to Office action of April 7, 2004

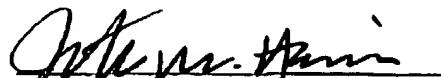
As explained previously, Passint teaches basing priority on header information and age, while Gotwald teaches basing priority on source address, destination address, the data type and the connection type. However, neither Passint nor Gotwald, nor a combination of Passint and Gotwald, teaches or suggests "a router" that "assigns the network ports higher priority than the microprocessor ports and prioritizes message packets associated with each of the network ports and each of the microprocessor ports according to a predetermined prioritization and according to a least-recently-granted prioritization" as required in claim 30. For at least these reasons, Applicants submit that claim 30 and all claims that depend from claim 30 are allowable.

III. CONCLUSIONS

In the course of the foregoing discussions, Applicants may have at times referred to claim limitations in shorthand fashion, or may have focused on a particular claim element. This discussion should not be interpreted to mean that the other limitations can be ignored or dismissed. The claims must be viewed as a whole, and each limitation of the claims must be considered when determining the patentability of the claims. Moreover, it should be understood that there may be other distinctions between the claims and the prior art which have yet to be raised, but which may be raised in the future.

Applicants respectfully request that a timely Notice of Allowance be issued in this case. If any fees or time extensions are inadvertently omitted or if any fees have been overpaid, please appropriately charge or credit those fees to Hewlett-Packard Company Deposit Account Number 08-2025 and enter any time extension(s) necessary to prevent this case from being abandoned.

Respectfully submitted,



Jonathan M. Harris
PTO Reg. No. 44,144
CONLEY ROSE, P.C.
(713) 238-8000 (Phone)
(713) 238-8008 (Fax)
ATTORNEY FOR APPLICANTS

HEWLETT-PACKARD COMPANY
Intellectual Property Administration
Legal Dept., M/S 35
P.O. Box 272400
Fort Collins, CO 80527-2400